

[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

**PORTAL**  
US Patent & Trademark Office

Search:  The ACM Digital Library  The Guide

## THE ACM DIGITAL LIBRARY

 [Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used [sampling](#) and [quicksort with buckets](#)

Found 1,297 of 123,929

Sort results by

Save results to a Binder

Try an [Advanced Search](#)

Display results

Search Tips

Try this search in [The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale 

### [1 A quantitative comparison of parallel computation models](#)

Ben H. H. Juurlink, Harry A. G. Wijshoff

August 1998 **ACM Transactions on Computer Systems (TOCS)**, Volume 16 Issue 3

Full text available:  [pdf\(1.06 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In recent years, a large number of parallel computation models have been proposed to replace the PRAM as the parallel computation model presented to the algorithm designer. Although mostly the theoretical justifications for these models are sound, and many algorithmic results were obtained through these models, little experimentation has been conducted to validate the effectiveness of these models for developing cost-effective algorithms and applications on existing hardware platforms. In ...

**Keywords:** parallel computation models, performance evaluation

### [2 Query evaluation techniques for large databases](#)

Goetz Graefe

June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2

Full text available:  [pdf\(9.37 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

**Keywords:** complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

### [3 A comparison of sorting algorithms for the connection machine CM-2](#)

Guy E. Blelloch, Charles E. Leiserson, Bruce M. Maggs, C. Greg Plaxton, Stephen J. Smith, Marco Zagha

June 1991 **Proceedings of the third annual ACM symposium on Parallel algorithms and architectures**

Full text available:  [pdf\(1.68 MB\)](#)

Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

### [4 Efficient low-contention parallel algorithms](#)

Phillip B. Gibbons, Yossi Matias, Vijaya Ramachandran

August 1994 **Proceedings of the sixth annual ACM symposium on Parallel algorithms and architectures**

Full text available:  pdf(1.39 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The queue-read, queue-write (QRQW) PRAM model [GMR94] permits concurrent reading and writing, but at a cost proportional to the number of readers/writers to a memory location in a given step. The QRQW model reflects the contention properties of most parallel machines more accurately than either the well-studied CRCW or EREW models: the CRCW model does not adequately penalize algorithms with high contention to shared memory locations, while the EREW model is too strict in its insistence on z ...

**5 Sorting**

W. A. Martin

December 1971 **ACM Computing Surveys (CSUR)**, Volume 3 Issue 4Full text available:  pdf(1.88 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The bibliography appearing at the end of this article lists 37 sorting algorithms and 100 books and papers on sorting published in the last 20 years. The basic ideas presented here have been abstracted from this body of work, and the best algorithms known are given as examples. As the algorithms are explained, references to related algorithms and mathematical or experimental analyses are given. Suggestions are then made for choosing the algorithm best suited to a given situation.

**6 High-performance sorting on networks of workstations**

Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, David E. Culler, Joseph M. Hellerstein, David A. Patterson

June 1997 **ACM SIGMOD Record , Proceedings of the 1997 ACM SIGMOD international conference on Management of data**, Volume 26 Issue 2Full text available:  pdf(1.53 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We report the performance of NOW-Sort, a collection of sorting implementations on a Network of Workstations (NOW). We find that parallel sorting on a NOW is competitive to sorting on the large-scale SMPs that have traditionally held the performance records. On a 64-node cluster, we sort 6.0 GB in just under one minute, while a 32-node cluster finishes the Datamation benchmark in 2.41 seconds. Our implementations can be applied to a variety of disk, memory, and processor configura ...

**7 Computing Curricula 1991**June 1991 **Communications of the ACM**, Volume 34 Issue 6Full text available:  pdf(2.87 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

ACM first published recommendations for undergraduate programs in computer science in 1968 in a report called "Curriculum '68." The report was produced as an activity of the ACM Education Board, which since then has been providing updates to recommendations for computer science programs as well as recommendations for other academic programs in computing.

**8 A Characterization of Ten Hidden-Surface Algorithms**

Evan E. Sutherland, Robert F. Sproull, Robert A. Schumacker

January 1974 **ACM Computing Surveys (CSUR)**, Volume 6 Issue 1Full text available:  pdf(4.47 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**9 Fast deterministic approximate and exact parallel sorting**

Torben Hagerup, Rajeev Raman

August 1993 **Proceedings of the fifth annual ACM symposium on Parallel algorithms and architectures**Full text available:  pdf(1.33 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**10 On randomization in sequential and distributed algorithms**

Rajiv Gupta, Scott A. Smolka, Shaji Bhaskar

March 1994 **ACM Computing Surveys (CSUR)**, Volume 26 Issue 1Full text available:  pdf(8.01 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Probabilistic, or randomized, algorithms are fast becoming as commonplace as conventional deterministic algorithms. This survey presents five techniques that have been widely used in the design of randomized algorithms. These techniques are illustrated using 12 randomized algorithms—both sequential and distributed—that span a wide range of applications, including primality testing (a classical problem in number theory), interactive probabilistic proofs ...

**Keywords:** Byzantine agreement, CSP, analysis of algorithms, computational complexity, dining philosophers problem, distributed algorithms, graph isomorphism, hashing, interactive probabilistic proof systems, leader election, message routing, nearest-neighbors problem, perfect hashing, primality testing, probabilistic techniques, randomized or probabilistic algorithms, randomized quicksort, sequential algorithms, transitive tournaments, universal hashing

**11 Network performance modeling for PVM clusters**

Mark J. Clement, Michael R. Steed, Phyllis E. Crandall

November 1996 **Proceedings of the 1996 ACM/IEEE conference on Supercomputing (CDROM)**

Full text available:  pdf(380.22 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

The advantages of workstation clusters as a parallel computing platform include a superior price-performance ratio, availability, scalability, and ease of incremental growth. However, the performance of traditional LAN technologies such as Ethernet and FDDI rings are insufficient for many parallel applications. This paper describes APACHE (Automated Pvm Application CCharacterization Environment), an automated analysis system that uses an application-independent model for predicting the impact ...

**12 Parallel construction of multidimensional binary search trees**

Ibraheem Al-Furaih, Srinivas Aluru, Sanjay Goil, Sanjay Ranka

January 1996 **Proceedings of the 10th international conference on Supercomputing**

Full text available:  pdf(1.09 MB)

Additional Information: [full citation](#), [references](#), [index terms](#)

**13 Further applications of random sampling to computational geometry**

K L Clarkson

November 1986 **Proceedings of the eighteenth annual ACM symposium on Theory of computing**

Full text available:  pdf(899.90 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

**14 Coarse grained parallel computing on heterogeneous systems**

Pat Morin

February 1998 **Proceedings of the 1998 ACM symposium on Applied Computing**

Full text available:  pdf(750.61 KB)

Additional Information: [full citation](#), [references](#), [index terms](#)

**Keywords:** bulk synchronous parallel, coarse grained multicomputer, heterogeneous computing, parallel algorithms

**15 Using attributed grammars to test designs and implementations**

A. G. Duncan, J. S. Hutchison

March 1981 **Proceedings of the 5th international conference on Software engineering**

Full text available:  pdf(789.78 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

We present a method for generating test cases that can be used throughout the entire life cycle of a program. This method uses attributed translation grammars to generate both inputs and outputs, which can then be used either as is, in order to test the specifications, or in conjunction with automatic test drivers to test an implementation against the specifications. The grammar can generate test cases either randomly or systematically. The attributes are used to guide the genera ...

**16 Query processing in main memory database management systems**

Tobin J. Lehman, Michael J. Carey

June 1986 **ACM SIGMOD Record , Proceedings of the 1986 ACM SIGMOD international conference on Management of data**, Volume 15 Issue 2Full text available:  pdf(1.43 MB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Most previous work in the area of main memory database systems has focused on the problem of developing query processing techniques that work well with a very large buffer pool. In this paper, we address query processing issues for memory resident relational databases, an environment with a very different set of costs and priorities. We present an architecture for a main memory DBMS, discussing the ways in which a memory resident database differs from a disk-based database. ...

**17 The algorithms course: teaching recent advances in the state of the art**

W. Douglas Maurer

March 2002 **The Journal of Computing in Small Colleges**, Volume 17 Issue 4Full text available:  pdf(84.86 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Many proofs of properties of algorithms may be simplified by the use of the method of indicator random variables. In addition, the relative speed of competing algorithms for the minimum-spanning-tree problem, and also for the problem of finding shortest paths in a graph, is affected by the development of a data structure called the Fibonacci heap. These are just the most important of the recent advances in the state of the art of analyzing algorithms which, in our opinion, need to be taught as ...

**18 Implementation and evaluation of an efficient parallel Delaunay triangulation algorithm**

Jonathan C. Hardwick

June 1997 **Proceedings of the ninth annual ACM symposium on Parallel algorithms and architectures**Full text available:  pdf(1.33 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**19 Developing a practical projection-based parallel Delaunay algorithm**

Guy E. Blelloch, Gary L. Miller, Dafna Talmor

May 1996 **Proceedings of the twelfth annual symposium on Computational geometry**Full text available:  pdf(1.17 MB)Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**20 An optimal algorithm for approximate nearest neighbor searching fixed dimensions**

Sunil Arya, David M. Mount, Nathan S. Netanyahu, Ruth Silverman, Angela Y. Wu

November 1998 **Journal of the ACM (JACM)**, Volume 45 Issue 6Full text available:  pdf(287.94 KB)Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Consider a set of  $S$  of  $n$  data points in real  $d$ -dimensional space,  $R_d$ , where distances are measured using any Minkowski metric. In nearest neighbor searching, we preprocess  $S$  into a data structure, so that given any query point  $q \in R_d$ , is the closest point of  $S$  to  $q$  can be reported quickly. Given any po ...

**Keywords:** approximation algorithms, box-decomposition trees, closet-point queries, nearest neighbor searching, post-office problem, priority search

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2003 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:

 [Adobe Acrobat](#) [QuickTime](#) [Windows Media Player](#) [Real Player](#)